Lecture #6: Finite-state Markov Chains: Transient Behavior

Randy Cogill
SYS 605 - Stochastic Systems
Fall 2007

Transient Behavior

- Before we looked at steady-state behavior of Markov Chains
- I.e., "How does a process behave in the long-term?"
- Now we look at **transient behavior** of Markov Chains
- I.e., "What happens along the way to a steady-state?"

Example: The game of Craps

- Recall the 'Pass line bet'
- Here's how it works:
  - Roll two dice. This is the 'come-out' roll
  - If total on dice is 7 or 11, we win
  - If total on dice is 2, 3 or 12, we lose
  - Otherwise, the total is called the 'point' and we continue to roll
  - Keep rolling until we either roll the point again or roll 7
  - If we roll the point, we win
  - If we roll 7, we lose
- Now ask, "How many rolls do we make per round?"

Hitting times

- Suppose we have a set $\mathcal{S} \subset \mathcal{X}$
- Starting in state $x \in \mathcal{X}$, time to reach $\mathcal{S}$ is the **hitting time**
- For a given sample path $X_0(\omega), X_1(\omega), \ldots$, hitting time is

$$T_{x,\mathcal{S}}(\omega) = \inf\{t \geq 0 \,|\, X_t(\omega) \in \mathcal{S},\ X_0(\omega) = x\}$$

- We want probabilities and expected values of hitting times

Probabilities of hitting times

- Consider the cost function defined by

$$r(x) = \begin{cases} 1 & \text{if } x \in \mathcal{S} \\ 0 & \text{otherwise} \end{cases}$$

- Create the 'stopped' process, where for each t

$$Y_t(\omega) = X_{\min\{t, T_{x,\mathcal{S}}\}}(\omega)$$

- The probability that $T_{x,\mathcal{S}} \leq k$ is

$$\begin{aligned} [\eta_k]_x &= \mathbf{P}(T_{x,\mathcal{S}} \leq k) \\ &= \mathbf{E}[r(Y_k) \,|\, Y_0 = x] \end{aligned}$$

Probabilities of hitting times (cont.)

- The stopped process is also a Markov chain

- Can create the transition matrix $\widehat{P}$ for the stopped MC:
  - For each $x \in \mathcal{S}$, replace column x of P with e_x
  - That is, make state x an absorbing state

- In terms of the transition matrix, $\mathbf{P}(T_{x,\mathcal{S}} \leq k)$ is

$$\eta_k = \left(\widehat{P}^T\right)^k r$$

Example: The game of Craps

- What is probability of completing the game in k or fewer rolls?
- Transition matrix is

$$\widehat{P} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 2/9 & 1 & 0 & 1/12 & 1/9 & 5/36 & 5/36 & 1/9 & 1/12 \\ 1/9 & 0 & 1 & 1/6 & 1/6 & 1/6 & 1/6 & 1/6 & 1/6 \\ 1/12 & 0 & 0 & 3/4 & 0 & 0 & 0 & 0 & 0 \\ 1/9 & 0 & 0 & 0 & 13/18 & 0 & 0 & 0 & 0 \\ 5/36 & 0 & 0 & 0 & 0 & 25/36 & 0 & 0 & 0 \\ 5/36 & 0 & 0 & 0 & 0 & 0 & 25/36 & 0 & 0 \\ 1/9 & 0 & 0 & 0 & 0 & 0 & 0 & 13/18 & 0 \\ 1/12 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 3/4 \end{bmatrix}$$

- We use $r = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}^T$

Example: Craps (cont.)

- Let $x =$ 'come out'
- As a function of k, $\mathbf{P}(T_{x,\mathcal{S}} \leq k)$ is below:

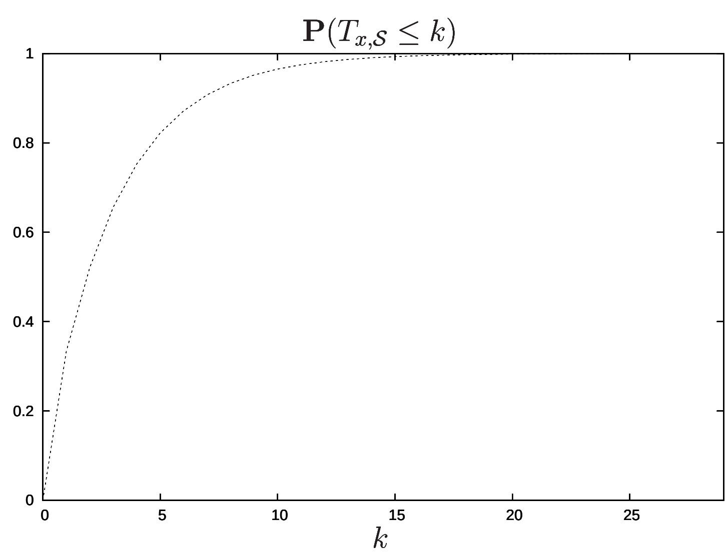

Total cost and hitting times

- Consider the cost function defined by

$$g(x) = \begin{cases} 0 & \text{if } x \in \mathcal{S} \\ 1 & \text{otherwise} \end{cases}$$

- Note that

$$T_{x,\mathcal{S}}(\omega) = \sum_{k=0}^{\infty} g(Y_k(\omega))$$

- For the stopped process Y_t, the expected hitting time from $x \in \mathcal{X}$ is

$$\mathbf{E}\left[\sum_{k=0}^{\infty} g(Y_t) \,\middle|\, Y_0 = x\right]$$

Total cost and hitting times (cont.)

- Can create the transition matrix $\widehat{P}$ for the stopped MC:
  - For each $x \in \mathcal{S}$, replace column x of P with e_x
- Denote the expected hitting from x by

$$[\gamma]_x = \mathbf{E}\left[\sum_{k=0}^{\infty} g(Y_t) \,\middle|\, Y_0 = x\right] = \sum_{k=0}^{\infty} \mathbf{E}[g(Y_t) \,|\, Y_0 = x]$$

- In terms of the transition matrix, $\gamma = \sum_{k=0}^{\infty} \left(\widehat{P}^T\right)^k g$

Finiteness of expected hitting times

- In general, we may not have $\displaystyle\sum_{k=0}^{\infty}\left(\widehat{P}^T\right)^k g < \infty$
- **Theorem:** Finite if for all $x \in \mathcal{X}$, there is some $y \in \mathcal{S}$ with $x \to y$
- **Proof:**
  - Arrange the states so that $\mathcal{S} = \{x_1, \ldots, x_m\}$
  - The matrix $\widehat{P}^T$ is of the form

$$\widehat{P}^T = \begin{bmatrix} I & 0 \\ P_{T1}^T & P_{T2}^T \end{bmatrix}$$

  - Continued...

Finiteness of exp. hitting times (cont.)

- **Proof (cont.):**
  - The powers $(\widehat{P}^T)^k$ are of the form

$$(\widehat{P}^T)^k = \begin{bmatrix} I & 0 \\ Q_k & \left(P_{T2}^T\right)^k \end{bmatrix}$$

  - So,

$$\sum_{k=0}^{n}\left(\widehat{P}^T\right)^k g = \sum_{k=0}^{n}\left(P_{T2}^T\right)^k \mathbf{1}$$

  - Now we'll show that for all x, y,

$$\sum_{n=0}^{\infty}[(P_{T2}^T)^n]_{xy} < \infty$$

Finiteness of exp. hitting times (cont.)

- **Proof (cont.):**
  - $y \notin \mathcal{S}$ is transient if $y \to z$ for some $z \in \mathcal{S}$
  - $\sum_{k=0}^{n-1} [P^k]_{yx}$ gives expected # of visits to y from x in n steps
  - $\sum_{k=0}^{n-1} [(P_{T2}^T)^k]_{xy} \leq \sum_{k=0}^{n-1} [(P_{T2}^T)^k]_{yy}$
  - Therefore,

$$\sum_{n=0}^{\infty} [(P_{T2}^T)^n]_{xy} \leq \sum_{n=0}^{\infty} [(P_{T2}^T)^n]_{yy} < \infty$$

Computing hitting times

- Assume $\mathcal{S}$ is reachable from all $x \in \mathcal{X}$, and hence $\sum_{k=0}^{\infty} (\widehat{P}^T)^k g < \infty$
- Then the expected hitting time satisfies $\gamma = g + \widehat{P}^T \gamma$
  - That is,

$$\sum_{k=0}^{n} (\widehat{P}^T)^k g = g + \widehat{P}^T \left(\sum_{k=0}^{n-1} (\widehat{P}^T)^k g \right)$$

  - Taking the limit on both sides gives $\gamma = g + \widehat{P}^T \gamma$
- However, $(I - \widehat{P}^T)\gamma = g$ doesn't have a unique solution...

Total cost and hitting times (cont.)

- We also know $[\gamma]_x = 0$ for all $x \in \mathcal{S}$
- Arrange the states so that $\mathcal{S} = \{x_1, \dots, x_m\}$
- Recall that the matrix $\widehat{P}$ is of the form

$$\widehat{P} = \begin{bmatrix} I & P_{T1} \\ 0 & P_{T2} \end{bmatrix}$$

- This gives

$$\begin{aligned}(I - \widehat{P}^T)\gamma &= \begin{bmatrix} 0 & 0 \\ -P_{T1}^T & I - P_{T2}^T \end{bmatrix} \begin{bmatrix} 0 \\ z \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{0} \\ \mathbf{1} \end{bmatrix}\end{aligned}$$

Total cost and hitting times (cont.)

- We can compute z by solving $(I - P_{T2}^T)z = 1$
- **Theorem:** If $\mathcal{S}$ is reachable from all $x \in \mathcal{X}$, $(I - P_{T2}^T)^{-1}$ exists
- **Proof:**
  - If P_{T2} is a submatrix of P corresponding to transient states,

$$\sum_{n=0}^{\infty} [P_{T2}^n]_{xy} < \infty$$

    for all x, y
  - Continued on next slide...

Total cost and hitting times (cont.)

- **Proof (cont.):**

  - So $\sum_{n=0}^{\infty}(P_{T2}^T)^n < \infty$, and

$$\lim_{n\to\infty}(I - P_{T2}^T)\sum_{k=0}^{n}(P_{T2}^T)^k = \lim_{n\to\infty}\left(I - (P_{T2}^T)^{n+1}\right)$$
$$= I$$

  - Therefore, $(I - P_{T2}^T)^{-1}$ exists and is equal to

$$(I - P_{T2}^T)^{-1} = \sum_{n=0}^{\infty}(P_{T2}^T)^n$$

Total cost and hitting times (cont.)

- This finally gives us a way of computing expected hitting times

- The unique γ with $[\gamma]_x = 0$ for all $x \in \mathcal{S}$ is

$$\gamma = \begin{bmatrix} \mathbf{0} \\ (I - P_{T2}^T)^{-1}\mathbf{1} \end{bmatrix}$$

Example: The game of Craps

- We want the expected number of rolls per round
- Transition matrix is

$$\widehat{P} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 2/9 & 1 & 0 & 1/12 & 1/9 & 5/36 & 5/36 & 1/9 & 1/12 \\ 1/9 & 0 & 1 & 1/6 & 1/6 & 1/6 & 1/6 & 1/6 & 1/6 \\ 1/12 & 0 & 0 & 3/4 & 0 & 0 & 0 & 0 & 0 \\ 1/9 & 0 & 0 & 0 & 13/18 & 0 & 0 & 0 & 0 \\ 5/36 & 0 & 0 & 0 & 0 & 25/36 & 0 & 0 & 0 \\ 5/36 & 0 & 0 & 0 & 0 & 0 & 25/36 & 0 & 0 \\ 1/9 & 0 & 0 & 0 & 0 & 0 & 0 & 13/18 & 0 \\ 1/12 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 3/4 \end{bmatrix}$$

- We use $g = \begin{bmatrix} 1 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix}^T$

Example: Craps (cont.)

- Solving for γ gives

$$\gamma = \begin{bmatrix} 3.376 & 0 & 0 & 4.0 & 3.6 & 3.273 & 3.273 & 3.6 & 4.0 \end{bmatrix}^T$$

- Therefore, expected number of rolls is $[\gamma]_1 \approx 3.376$